

THE SCALING LIMIT OF RANDOM OUTERPLANAR MAPS

Alessandra Caraceni*

Abstract

A planar map is outerplanar if all its vertices belong to the same face. We show that random uniform outerplanar maps with n vertices suitably rescaled by a factor $1/\sqrt{n}$ converge in the Gromov-Hausdorff sense to $7\sqrt{2}/9$ times Aldous' Brownian tree. The proof uses the bijection of Bonichon, Gavoille and Hanusse [7].

1 Introduction

Since the early 90's a lot of effort has gone into the study of scaling limits for large combinatorial structures. The most emblematic result is, of course, the construction of the continuum random tree (CRT) by Aldous [3, 4, 5] as the scaling limit of various classes of oriented trees. The CRT has since been shown to be a universal object: see e.g. [1, 2, 10, 12, 14] and references therein, as well as the recent work [19]. In this work we shall establish that the CRT is also the scaling limit of uniform random large outerplanar maps.

Recall that a planar map is a proper embedding of a finite connected graph into the plane (or the sphere), considered up to continuous deformations. A recent breakthrough was achieved by Le Gall and Miermont [15, 18], who showed that several classes of random planar maps admit the so-called Brownian map as scaling limit. It has been observed, however, that – for some particular regimes – random planar maps with a unique macroscopic large face have a tree-like structure and admit the CRT as scaling limit; see [6, 8, 13, 19]. Our main result consists in a confirmation of this phenomenon for the case of outerplanar maps.

A map is *outerplanar* if all of its vertices are adjacent to the same face, which is dubbed the *outerface* and usually drawn as the infinite face in a planar embedding. Outerplanar maps constitute a well-studied combinatorial structure; in particular, they have a simple characterisation in terms of minors (a graph is outerplanar if and only if it does not contain $K_{2,3}$ nor K_4 as a minor [9]). See [20] for more characterisations of outerplanar graphs. In this work we shall restrict ourself to *simple* outerplanar maps, with no loops or multiple edges. As usual all of the maps considered here are *rooted*,

*Scuola Normale Superiore di Pisa, e-mail: alessandra.caraceni@sns.it




Figure 1: A rooted, simple outerplanar map.

that is endowed with a distinguished oriented edge such that the outerface is lying on its left. The tail of the root edge will be called the *root vertex*. Our main result is the following:

Theorem 1.1. *Let \mathbf{M}_n be a random uniform rooted simple outerplanar map with n vertices, and denote by d_{gr} the graph distance on the set of its vertices $V(\mathbf{M}_n)$. We have the following convergence in distribution for the Gromov-Hausdorff topology:*

$$\left(V(\mathbf{M}_n), \frac{d_{gr}}{\sqrt{n}} \right) \xrightarrow[n \rightarrow \infty]{(d)} \frac{7\sqrt{2}}{9} \cdot (\mathcal{T}_e, d),$$

where (\mathcal{T}_e, d) is the Brownian CRT of Aldous. We adopt here the normalisation of Le Gall [14] by considering \mathcal{T}_e as constructed from a normalised Brownian excursion.

The first ingredient in our proof is a way to relate outerplanar maps to plane trees; this will be done using the bijection of Bonichon, Gavoille and Hanusse [7] between the set of (simple and rooted) outerplanar maps with n vertices and a special class of bicoloured plane trees with n vertices which is described in Section 2. The plan of the proof then partially follows that of [10], in which Curien, Haas and Kortchemski prove the convergence of random dissections to a scalar multiple of the CRT. More specifically, we will show that the distances on an outerplanar map are roughly proportional to the distances on the associated tree. To this end, we describe throughout Section 4 an

algorithm that, given a bicoloured tree and a vertex v , yields the length of a geodesic path from v to the root vertex in the associated outerplanar map. When applied to the model of a bicoloured Galton-Watson tree conditioned to survive (presented in Section 5) this algorithm yields a Markov chain whose mean increment (under the stationary distribution) gives the asymptotical proportionality constant between the metric on a large outerplanar map and that of its associated tree. The distances between arbitrary pairs of points are finally controlled by a large deviations estimate, see Sections 6 and 7.

2 Outerplanar maps and plane trees

As mentioned earlier, the first ingredient needed for our discussion is a bijection found in [7], which enables the coding of outerplanar maps as bicoloured trees of a certain class. More specifically,

Definition 2.1. We say that a rooted plane tree τ is *bicoloured* if each of its vertices is coloured either black or white; we shall say that τ is *well bicoloured* if it is bicoloured and all of the vertices in its rightmost branch are coloured white (see Figure 2a). We shall henceforth simply write *well bicoloured tree* when referring to a well bicoloured rooted plane tree.




Figure 2: Figure (a) shows a well bicoloured tree. Vertices u and v are *unrelated*. Vertex w is the *target* of v . Figure (b) represents its image via the map Ψ .

Some working knowledge of the explicit bijection is needed in the sections that follow, and thus part of the construction is included for future reference. We adopt notation coherent with that of [7]: given two distinct vertices in a plane tree, we call them *unrelated* if neither is an ancestor of the other.

Let now τ be a well bicoloured tree. For each black vertex v of τ define the *target* of v to be its next unrelated vertex in a clockwise contour of τ . Define $\Psi(\tau)$ as the rooted outerplanar map obtained by joining each black vertex of τ to its target (via an edge that leaves the rightmost corner of the black vertex and enters the target from the leftmost

corner available); root the map on the edge joining the former root of the tree to its leftmost child, oriented in such a way that the former root is the tail, and forget the colouring of τ (see Figure 2b for an example).

Notice that the order in which the additional edges are drawn is not relevant, and that the root edge will have the infinite face on its left side. An inverse of Ψ can be constructed explicitly, but we refer the reader to [7] for the details and proof, since all that we will need is the following:

Theorem 2.1 (Bonichon, Gavoille, Hanusse [7]). *The map Ψ is a bijection between well bicoloured trees with n vertices and simple rooted outerplanar maps with n vertices.*

Such a bijection, together with the fact that the scaling limit for plane trees is the CRT, constitutes the basis for our future discussion; notice that, however, it is not at all clear how the colouring affects the metric in the switch from tree to map: distances between corresponding vertices are, in general, smaller when computed on the map (if two vertices are adjacent in the tree then they are in the corresponding map, but not vice-versa). This means we cannot easily employ the result for plane trees to make deductions on outerplanar maps. Most of the following sections will develop ways to control the outerplanar map metric via easily readable information about its corresponding tree.

3 Rough localisation of geodesics

We delve now into the central problem of the rather unclear relationship between distances on an outerplanar map and distances on its corresponding plane tree: given a well bicoloured tree τ , we wish to compute distances on the map $\Psi(\tau)$. We restrict ourselves, in this section and many of the subsequent ones, to distances from the root vertex; in Section 7, before the proof of Theorem 1.1, we will give a way to bound distances between arbitrary vertices of an outerplanar map with a function of distances from the root. In what follows, geodesics to the root in $\Psi(\tau)$ are built in a step-by-step manner, using local information about the tree structure of τ and its colouring.

We shall refer to geodesics of $\Psi(\tau)$ as map-geodesics; since τ and $\Psi(\tau)$ have the same vertex set, any path in $\Psi(\tau)$ can be interpreted as a sequence of vertices $v_0 \dots v_n$ of τ , where for each i between 0 and $n - 1$ vertices v_i and v_{i+1} are either neighbours in τ (parent and child, in any order), or a black vertex and its target (again, a priori, in any order).

We need some additional notation: for each vertex u of τ , if u is not the root, we write $p(u)$ for its parent; if u has children in τ , we call $r(u)$ its rightmost child; finally, if u is a black vertex, we write $t(u)$ for its target. Also, given three vertices u , v and w of τ , we say that (v, w) is a *separating pair* for u (from the root of the tree) if v is a strict ancestor of u and w is a child of v lying to the right of u ; see Figure 3.

We write $d_M(u, v)$ for the graph distance between vertices u and v in the map $\Psi(\tau)$, and simply write $d_M(u)$ for $d_M(u, \emptyset)$, where \emptyset is the root of τ . We are ready to prove the following:





Figure 3: (v, w) is a separating pair for u (also, it is a separating pair for $r(u)$, $t(u)$, $p(u)$, or indeed for any of the circled vertices).

Proposition 3.1. *In a well bicoloured tree τ , let (v, w) be a separating pair for u . Then any map-geodesic from u to the root of τ will pass through v or w (possibly both). Consequently, any map-geodesic from u to the root can be constructed by starting from u and iteratively applying one of the maps t , r or p , so that r or p are applied to white vertices, and p or t are applied to black vertices.*

Proof. Let S be the set of (strict) descendants of v that lie strictly to the left of w , and $u_0 u_1 \dots u_n$ a map-geodesic such that $u = u_0$ and $u_n = \emptyset$ is the root of the tree. Clearly, $u \in S$ and $u_n \notin S$; take the minimum i such that $u_i \notin S$, and consider its relation to u_{i-1} . Children of elements in S are in S , and for any pair $(x, t(x))$ in the tree (where x is a black vertex), either both vertices belong to S or neither does, except for the case where $x \in S$ and $t(x) = w$. Hence either u_i is the parent of u_{i-1} (therefore $u_i = v$) or u_i is the target of u_{i-1} , which can only be the case if $u_i = w$, and thus is the first part of the proposition established.

Let us consider what this implies in term of distances: if (v, w) is a separating pair for u , then $d_M(u) > d_M(v)$ or $d_M(u) > d_M(w)$; that is, since $|d_M(v) - d_M(w)| \leq 1$, $d_M(u) \geq \max\{d_M(v), d_M(w)\}$. Now, consider a map-geodesic from u to the root. If x is a child of u distinct from $r(u)$, then such a geodesic does not go through x , because $(u, r(u))$ is a separating pair for x , and thus $d_M(x) \geq d_M(u)$ (whereas, if there were a geodesic from u to the root that involved x , we would have $d_M(x) < d_M(u)$). If u is the target of some y , then consider $p(u)$, which must be an ancestor of y : $(p(u), u)$ is a separating pair for y , and thus $d_M(y) \geq d_M(u)$, so a map-geodesic from u to the root does not go through y . Only three possibilities remain: either the geodesic moves from u to $p(u)$, or to $t(u)$, or to $r(u)$.

Clearly, a white vertex u in a map-geodesic to the root will be followed by either



$r(u)$ or $p(u)$, since it is not directly connected to any target in the map. Suppose, on the contrary, that u is a black vertex; it will be followed in a map-geodesic to the root by $p(u)$ or $t(u)$, never by a child: this is because $(p(t(u)), t(u))$ is a separating pair for u ; if the map-geodesic does not move from u to $t(u)$, then it passes through $p(t(u))$, which has map-distance at most 2 from u , and at least 2 from any child of u (it cannot be the target of one, since children of u have $t(u)$ or other children of u as targets, and is not connected to strict descendants of u in the tree); as a consequence, the map-geodesic does not go through $r(u)$. \square


4 An algorithm to compute map-distances

Suppose we have a well bicoloured tree τ and a vertex x of τ of height n ; we know that $d_M(x)$, the map-distance between x and the root of τ , is no more than n . From now on, we write $d(\tau, x)$ for the map distance $d_M(x)$ between x and the root of τ . We propose to compute $d(\tau, x)$ via a recursive algorithm which takes the pair (τ, x) as input and outputs a pair (τ', x') , where x' is a vertex of τ' such that $d(\tau', x') = d(\tau, x) - 1$; this way the number of iterations needed for the algorithm to output a well bicoloured tree and its root is exactly the length of a geodesic path from x to the root of τ .


Given (τ, x) , consider the path *in the tree* leading from x to the root; thanks to Proposition 3.1 we know that a map-geodesic from x to the root cannot involve any of the vertices that lie strictly to the left of this path (parents, targets and rightmost children of vertices that are part of the path or lie to the right of it cannot lie to its left). We may thus safely erase all such vertices from τ , and we will always output pairs (τ', x') such that no vertices lie strictly to the left of the tree path leading from x' to the root of τ' .

In what follows, given a tree τ and a vertex x , we will write τ^x for the subtree of τ formed by x and its descendants; given a tree τ and a subtree $\tilde{\tau}$ (which does not include the root of τ), we will write $\tau \setminus \tilde{\tau}$ for the rooted plane tree obtained from τ by erasing all vertices of $\tilde{\tau}$ and any edges adjacent to those vertices.


The description of the algorithm follows.




[w₀] Suppose x is a white leaf of τ ; then a map-geodesic to the root necessarily moves from x to $p(x)$ (hence $d(\tau, p(x)) = d(\tau, x) - 1$). Consider the tree $\tau' = \tau \setminus x$ and the pair $(\tau', p(x))$ (where we still write $p(x)$ for the obvious image of the original vertex of τ in τ'); it is clear that $d(\tau', p(x)) = d(\tau, p(x))$ (because the only vertex removed is x , which was further from the root than $p(x)$) and so it equals $d(\tau, x) - 1$.



[w_{>0}] If x is a white vertex of τ and has offspring, the matter is more complicated. Proposition 3.1 ensures that a geodesic to the root moves to either $r(x)$ or $p(x)$, but it is not clear which: we need to distinguish two cases.




[w_{>0}.1] Suppose that $x \neq r(p(x))$, that is x has some right siblings; then $(p(x), r(p(x)))$ is a separating pair for $r(x)$, and thus $d(\tau, p(x)) \leq d(\tau, r(x))$; hence there is a geodesic moving from x to $p(x)$, and $d(\tau, p(x)) = d(\tau, x) - 1$. We choose to follow such a geodesic and define τ' to be $\tau \setminus \tau^x$, and output $(\tau', p(x))$ so that, as before, we have $d(\tau', p(x)) = d(\tau, x) - 1$.




[w_{>0}.2] Suppose now that $x = r(p(x))$. We build the new tree τ' by erasing τ^x from the original tree and rerooting the subtree $\tau^{r(x)}$ onto $p(x)$, by identifying $p(x)$ with $r(x)$, thus merging them into a single vertex y ; the colour of y is set to white if and only if both $p(x)$ and $t(x)$ were white in τ . We need to show that $d(\tau', y) = d(\tau, x) - 1$: there is an obvious map sending vertices of τ that are not in $\tau^x \setminus \tau^{r(x)}$ to vertices of τ' ; the map is 1-on-1 with the exceptions of $p(x)$ and $r(x)$, which are both sent to y . Neighbours in τ are sent to neighbours in τ' , and (since $x = r(p(x))$ in τ) the target of a vertex in τ becomes the target of its image in τ' : hence $d(\tau', y) \leq \min\{d(\tau, p(x)), d(\tau, r(x))\}$. On the other hand, any map-path in τ' can be lifted to a map-path in τ (by appropriately selecting a pre-image for y as first step of the path), which gives equality.




[b] If x is a black vertex of τ , then the new option of *jumping* to $t(x)$ presents itself. As by Proposition 3.1, a geodesic will either move to $p(x)$ or to $t(x)$. We need to deal with three separate cases:



[b.1] Suppose x , $t(x)$ and $r(p(x))$ are distinct: that is to say, x has at least two right siblings; in this case, $(p(x), r(p(x)))$ is a separating pair for $t(x)$, hence there is a geodesic moving from x to $p(x)$; we thus set τ' to be $\tau \setminus \tau^x$, and output $(\tau', p(x))$.



[b.2] Suppose now that x has only one right sibling, which is therefore $t(x)$ as well as $r(p(x))$. We build the new tree τ' by simply erasing τ^x and identifying vertices $p(x)$ and $t(x)$, merging them into a single vertex y to be coloured white if and only if both of the original vertices were white in τ , and output (τ', y) . We have $d(\tau', y) = d(\tau, x) - 1$ by roughly the same argument as previously.



[b.3] The last case is that of x being the rightmost child of its parent; in this case $(p(t(x)), t(x))$ is a separating pair for $p(x)$, so $d(\tau, t(x)) \leq d(\tau, p(x))$, and we may assume a geodesic to the root does jump from x to its target. We can thus build τ' by erasing *all that lies left of* $t(x)$, and output $(\tau', t(x))$.

Notice that, in all cases listed except for the very last one, the output vertex x' has height $|x| - 1$ in τ' , whereas the jump made in the last case (the one marked [b.3]) may

lead to a vertex x' of arbitrarily smaller height. Also, the information on τ that the algorithm uses to select the appropriate τ' is entirely local (child structure of x and its parent) with the exception of the last case, which requires to make changes to parts of τ that are, a priori, arbitrarily far from x .

In the spirit of making each step by the algorithm entirely determined by local information, which will in turn entail precious independence properties as soon as we switch to a random setting, we add extra data to inputs and outputs: we let the algorithm run on triples of the form (τ, x, s) , where s is one of four *states*, and output a triple (τ', x', s') .


Three of the states simply mimic the cases listed above: we call them w_0 , $w_{>0}$ and b ; a fourth state, labelled j for *jump*, is devised to deal specifically with situations that fall under case [b.3]: the idea is that, instead of simply outputting an entirely different tree paired with the target of the jump, the algorithm goes into a *jump* state; it proceeds modifying the tree a little at a time until it reaches the original target, at which point it 'lands' in one of the three non-jump states.

We propose to re-define outputs according to the state of the input triple, keeping in mind that they mostly adhere to the preceding description for pairs; if the output state in the triple (τ', x', s') is known *not* to be j , then it is determined by (τ', x') . The initial state, in particular, is not jump: given τ and a vertex x , it can be determined as being


- w_0 , if x is a white leaf of τ ;
- $w_{>0}$, if x is white and it is not a leaf in τ ;
- b , if x is black.

The behaviour of the algorithm for an input triple of the form (τ, x, w_0) or $(\tau, x, w_{>0})$, and for input triples (τ, x, b) with x having one or more right siblings, is exactly that described in cases $[w_0]$, $[w_{>0}.1]$, $[w_{>0}.2]$, $[b.2]$ and $[b.1]$. Namely, the tree τ' and a vertex x' are produced, and the state s' is selected again among the three non-jump states (w_0 , $w_{>0}$, b) according to the colour and degree of the output vertex x' (white leaf, white non-leaf, black).

We now describe the behaviour of the algorithm when the input is of the form (τ, x, b) , and x is the rightmost child of its parent. As explained earlier (case [b.3]), the next vertex in a map-geodesic to the root would be (without loss of generality) $t(x)$. We distinguish yet two subcases.




[b.3.1] If $t(x)$ has height $|x| - 1$ (that is, if $p(x)$ has a right sibling: again, a local property) then define τ' by erasing $\tau^{p(x)}$ from τ and output $(\tau', t(x), s')$, where s' is – again – one of w_0 , $w_{>0}$, b , according to properties of $t(x)$ in τ' .




[b.3.2] If, however, $p(x)$ has no right siblings, then it is time to finally put the jump state j to use. We define τ' to be $\tau \setminus \tau^x$ and output $(\tau', p(x), j)$. Notice that now $p(x)$ has height $|x| - 1$ in τ' , and that the identity of $t(x)$ can still be recovered in τ' (even though vertex x has been erased) by the sole knowledge of $p(x)$: if $p(x)$ were black, $t(x)$ would be the target of $p(x)$.

We finally give instructions for the algorithm to follow when confronted with a jump state. Suppose we have an input (τ, x, j) ;



[j.1] if $p(x)$ has no right siblings, then output $(\tau \setminus \tau^x, p(x), j)$; this way, the vertex on which the geodesic should land is still the target of $p(x)$ (or would be if $p(x)$ were black) and its map-distance from the root remains unchanged;



[j.2] if $p(x)$ has right siblings, then the leftmost one (call it x') is the (image of the) vertex the geodesic was supposed to land on; output $(\tau \setminus \tau^{p(x)}, x', s')$, with s' being appropriately chosen among the three non-jump states.

We now summarise the key properties of the algorithm as just described: given a well bicoloured tree τ and a vertex x of height n , we can generate a sequence of $n + 1$ triples (τ_i, x_i, s_i) , with i ranging from 0 to n , where

- for each i , τ_i is a well bicoloured tree, x_i is a vertex of τ and s_i is one of four states $(w_0, w_{>0}, b, j)$;
- τ_0 is the tree τ deprived of all that lies left of the (tree) path from the root to x , and $x_0 = x$;
- $(\tau_{i+1}, x_{i+1}, s_{i+1})$ is obtained from (τ_i, x_i, s_i) as described, with changes of a ‘local’ nature;
- for each i , the height of x_i in τ_i is $n - i$; consequently, x_n is the root of τ_n ;
- for each i between 0 and $n - 1$ such that $s_i \neq j$, $d(\tau_{i+1}, x_{i+1}) = d(\tau_i, x_i) - 1$; therefore,

$$d(\tau, x) = \sum_{0 \leq i < n | s_i \neq j} 1 = n - |\{0 \leq i < n | s_i = j\}|; \quad (\star)$$

that is, the distance $d(\tau, x)$ is the number of non-jump states appearing in the input triples on which the algorithm runs (indeed, it is also the number of non-jump states appearing in the n *output* triples, because s_n cannot be j and neither can s_0). We will make frequent use of this fact in what follows.

The time has come to run our algorithm on a random tree; in order to do this, an especially useful tool is the standard construction of the *geometric Galton-Watson tree*




Figure 4: Figure (a) represents the critical geometric Galton-Watson tree conditioned to survive; trees $\rho^1, \dots, \rho^n, \dots$ and $\lambda^1, \dots, \lambda^n, \dots$ are independent geometric Galton-Watson trees. Figure (b) represents $\text{cut}_2(T_\infty)$.

conditioned to survive, which will provide us with a way of unifying results given by the algorithm for vertices of arbitrary height into a single random variable. The next section is an introduction to this tool and to some of the notation needed for further progress.

5 The Galton-Watson tree conditioned to survive

In this section we briefly introduce the critical Galton-Watson tree conditioned to survive for a geometric offspring distribution; for a more general definition and further details see Section 12 of [16].

We build a random infinite tree, called T_∞ , in the following way: consider an infinite path $v_0 v_1 \dots v_n \dots$, called the *spine*; let $(\lambda^i)_{i \geq 0}$ be a sequence of independent critical geometric Galton-Watson trees (that is, with offspring distribution μ , where $\mu(k) = 2^{-k-1}$ for all natural numbers k), and let $(\rho^i)_{i \geq 0}$ be another such sequence, independent of the first; we consider the Galton-Watson trees as random rooted plane trees and, for each $i \geq 0$, attach λ^i to the spine by identifying its root and v_i , so that λ^i lies to the left of the spine; similarly, attach each ρ^i so that its root is identified with v_i and all of its vertices lie to the right of the spine (see Figure 4a); finally, root the random infinite tree thus obtained in v_0 .

The relevance of critical geometric Galton-Watson trees lies in the fact that, if θ is one such tree and θ_n is the random tree obtained by conditioning θ on having exactly n vertices, then θ_n is uniformly distributed over plane trees with n vertices; as a result, the tree T_∞ itself has much to do with random plane trees, as we shall now see.

We write $\text{cut}_n(T_\infty)$ for the pair (τ, v_n) , where τ is the (finite plane rooted) tree obtained from T_∞ by erasing every descendant of v_n which lies *strictly to the left* of v_{n+1} , together with v_{n+1} itself and all of its descendants: that is, τ consists of the path $v_0 \dots v_n$, with trees ρ^0, \dots, ρ^n attached to the right and trees $\lambda^0, \dots, \lambda^{n-1}$ attached to the left. Then we have the following standard result.

Lemma 5.1. *Let F be a non-negative real valued function defined on pairs (τ, u) , where τ is a finite rooted plane tree and u is a vertex in τ ; let θ be a Galton-Watson tree with critical geometric offspring distribution. Then for all $n \geq 0$*

$$\mathbb{P}[F(\text{cut}_n(T_\infty))] = \mathbb{P}\left[\sum_{\substack{u \in \theta \\ |u|=n}} F(\theta, u)\right].$$

Proof. For the left hand side we have

$$\begin{aligned} \mathbb{P}[F(\text{cut}_n(T_\infty))] &= \sum_{(\tau, u)} F(\tau, u) \mathbb{P}(\text{cut}_n(T_\infty) = (\tau, u)) = \\ &= \sum_{\tau} \sum_{\substack{u \in \tau \\ |u|=n}} F(\tau, u) \mathbb{P}(\text{cut}_n(T_\infty) = (\tau, u)) \end{aligned}$$

where (τ, u) ranges among all pairs formed by a finite (rooted plane) tree and a vertex u of height n in the tree.

Thus it is enough to show that, for all such pairs (τ, u) ,

$$\mathbb{P}(\text{cut}_n(T_\infty) = (\tau, u)) = \mathbb{P}(\theta = \tau).$$

But the probability $\mathbb{P}(\text{cut}_n(T_\infty) = (\tau, u))$ is easy to compute: consider the set formed by u and its ancestors in τ ; order them according to height, and label them u_0, \dots, u_n , so that $|u_i| = i$ (u_0 is the root and u_n is u). For $i = 0, \dots, n-1$, let τ_l^i be the subtree of τ formed by u_i and its descendants lying strictly to the left of u_{i+1} . Similarly, let τ_r^i be the subtree of τ formed by u_i and its descendants lying strictly to the right of u_{i+1} . Then

$$\begin{aligned} \mathbb{P}(\text{cut}_n(T_\infty) = (\tau, u)) &= \left(\prod_{i=0}^{n-1} \mathbb{P}(\lambda^i = \tau_l^i) \mathbb{P}(\rho^i = \tau_r^i) \right) \cdot \mathbb{P}(\rho^n = \tau^u) = \\ &= \left(\prod_{i=0}^{n-1} \mathbb{P}(\theta = \tau_l^i) \mathbb{P}(\theta = \tau_r^i) \right) \cdot \mathbb{P}(\theta = \tau^u). \end{aligned}$$

Now, for each vertex x in τ , let $c(x)$ be the number of children of x in τ ; clearly,

$$\mathbb{P}(\theta = \tau) = \prod_{x \in \tau} 2^{-c(x)-1}$$

by definition of θ .

On the other hand, for each vertex in tree τ_l^i (i between 1 and $n-1$) call $c_l(x)$ the number of its children in τ_l^i ; similarly, call $c_r(x)$ the number of children of x in τ_r^i , if x belongs to such a tree. Then

$$\mathbb{P}(\text{cut}_n(T_\infty) = (\tau, u)) = \prod_{i=0}^{n-1} \left\{ \prod_{x \in \tau_l^i} 2^{-c_l(x)-1} \prod_{x \in \tau_r^i} 2^{-c_r(x)-1} \right\} \cdot \prod_{x \in \tau^u} 2^{-c(x)-1}.$$

Consider any vertex x of τ such that $x \notin \{u_0, \dots, u_{n-1}\}$; then x appears only once in the expression above, as part of some tree τ_l^j or τ_r^j , or possibly of τ^u . Furthermore, the number of children of x in its subtree ($c_l(x)$ or $c_r(x)$ or $c(x)$) is exactly the same as the number of children $c(x)$ that x has in τ . Now, for $i = 0, \dots, n-1$, consider u_i ; it appears inside two of the products, as part of tree τ_l^i and tree τ_r^i , therefore it contributes to $\mathbb{P}(\text{cut}_n(T_\infty) = (\tau, u))$ with a factor $2^{-c_l(u_i)-1}2^{-c_r(u_i)-1}$; on the other hand, u_i has $c(u_i) = c_l(u_i) + c_r(u_i) + 1$ children in τ , so that its contribution to $\mathbb{P}(\theta = \tau)$ is a factor $2^{-c(u_i)-1} = 2^{-c_l(u_i)-c_r(u_i)-2}$.

Consequently, we have $\mathbb{P}(\text{cut}_n(T_\infty) = (\tau, u)) = \mathbb{P}(\theta = \tau)$, as wanted. \square

In order to adapt the notion of the critical geometric Galton-Watson tree conditioned to survive to our prior setting, we need to endow it with a random bicolouring. We do this by simply choosing the colour for each vertex of T_∞ uniformly at random with probability $1/2$. In what follows, since the context will determine whether or not trees are bicoloured, we will still write T_∞ for the object just introduced, namely the *critical geometric Galton-Watson tree conditioned to survive, uniformly bicoloured*.

Similarly, if τ is a (random) finite tree, it can be uniformly bicoloured by choosing a colour for each of its vertices independently and uniformly at random, conditionally on τ itself. This applies in particular when τ is a critical geometric Galton-Watson tree.

Finally, we remark that a bicoloured critical geometric Galton-Watson tree is not necessarily *well* bicoloured; given a finite bicoloured tree τ , we define the new tree τ° as the one obtained by adding a white leaf as rightmost child of the root of τ and recolouring the root white, so that τ° is well bicoloured; we will always think of τ as being embedded in τ° in the obvious way. One may of course consider map-distances on τ° : given a vertex u in τ , we write $d^\circ(\tau, u) = d(\tau^\circ, u)$ for the map distance between vertex u and the root in τ° . Analogously, we may consider $d^\circ(\text{cut}_n(T_\infty))$, which, if $\text{cut}_n(T_\infty) = (\tau, u)$, we take to mean $d(\tau^\circ, u)$.

6 The algorithm running on the infinite bicoloured tree

Let T_∞ be the critical geometric Galton-Watson tree conditioned to survive, uniformly bicoloured; call v_0, \dots, v_n, \dots the vertices on its spine (v_0 being the root). The distance $d^\circ(\text{cut}_n(T_\infty))$ is, for each positive integer n , a random variable which we wish to estimate (at least asymptotically in n).

We have described all through Section 4 an algorithm that can now be started on (τ°, x, s) , where $(\tau, x) = \text{cut}_n(T_\infty)$ and s depends on the colour and offspring of x in τ (therefore on T_∞ and n); this will yield a sequence of n states s_0, \dots, s_{n-1} , each a random variable taking values in the space of states $\Sigma = \{w_0, w_{>0}, b, j\}$, such that

$$d^\circ(\text{cut}_n(T_\infty)) = n - |\{0 \leq i < n \mid s_i = j\}|.$$

This sequence $(s_i)_{0 \leq i < n}$ is ‘almost’ a Markov Chain, in the sense clarified by the following fundamental proposition:

Proposition 6.1. Fix $n > 0$; take the random infinite tree T_∞ and consider the sequence s_0, \dots, s_{n-2} of the first $n - 1$ inputs for the algorithm from Section 4, started on (τ°, x, s_0) , where $(\tau, x) = \text{cut}_n(T_\infty)$. Then such a sequence has the same law as (the first $n - 1$ steps of) a Markov chain $(X_i)_{i \geq 0}$ with transition matrix

$$M = \begin{pmatrix} \uparrow & w_0 & w_{>0} & b & j \\ w_0 & 1/4 & 1/4 & 1/2 & 0 \\ w_{>0} & 1/16 & 5/16 & 5/8 & 0 \\ b & 3/32 & 7/32 & 7/16 & 1/4 \\ j & 1/8 & 1/8 & 1/4 & 1/2 \end{pmatrix}$$

and a random initial state distributed as $(1/4, 1/4, 1/2, 0)$.

This proposition plays a key role in finally establishing Theorem 1.1; it is, in fact, the motivation that led to the algorithm as described in Section 4, and its proof is nothing but a careful observation of how the various steps of the algorithm interact with the random element: in particular, some key independence properties are always preserved, mainly thanks to the Galton-Watson structure of subtrees and the nature of the geometric law.

In order to prove Proposition 6.1, we highlight those exact independence properties in a separate Lemma, for which some additional notation is needed. Suppose the algorithm is started on a triple (τ_0, x_0, s_0) , where x_0 has height n in τ_0 ; we write p^j for the j -th iteration of the parent function, so that $p^0(x_0) = x_0$ and $p^j(x_0)$ has height $n - j$ (in particular, $p^n(x_0)$ is the root of τ_0). We call ρ_0^j the subtree of τ_0 consisting of $p^{n-j}(x_0)$ and its descendants lying *strictly to the right* of $p^{n-j-1}(x_0)$ (so that the tree ρ_0^j is rooted in a vertex of height j in τ_0); ρ_0^n is simply $\tau_0^{x_0}$.

We repeat the same construction for all subsequent triples (τ_i, x_i, s_i) : x_i has height $n - i$; we write ρ_i^j for the subtree of τ_i consisting of $p^{n-i-j}(x_i)$ and its descendants lying *strictly to the right* of $p^{n-i-j-1}(x_i)$ (see Figure 5); as before, the root of ρ_i^j has height j in τ_i , and $\rho_i^{n-i} = \tau_i^{x_i}$.

We wish to prove the following:

Lemma 6.2. Fix $n > 0$ and consider the (random) triples (τ_i, x_i, s_i) obtained from T_∞ through the algorithm for i between 0 and $n - 2$. Then for each i ,

- if $s_i = w_0$ or $s_i = b$, then ρ_i^j , for $j = 1, \dots, n - i - 1$, is a sequence of independent uniformly bicoloured Galton-Watson trees;
- if $s_i = w_{>0}$, then ρ_i^j , for $j = 1, \dots, n - i - 1$, is a sequence of independent uniformly bicoloured Galton-Watson trees; also, $\tau_i^{r(x_i)}$ is a uniformly bicoloured Galton-Watson tree independent of the block of the ρ_i^j 's;




Figure 5: Structure of τ_i .

- if $s_i = j$, then ρ_i^j , for $j = 1, \dots, n-i-2$, is a sequence of independent uniformly bicoloured Galton-Watson trees, whilst $\rho_i^{n-i-1} = \{p(x_i)\}$ and $\rho_i^{n-i} = \{x_i\}$.

Proof. We proceed by induction on i ; for $i = 0$, all assertions are trivial by definition of T_∞ (more precisely, by the random structure of $\text{cut}_n(T_\infty)$): with the exception of ρ_0^0 , whose root is recoloured as white and has a white leaf attached as a rightmost child, $\rho_0^j = \rho^j$: all right trees attached to the spine of T_∞ are uniformly bicoloured Galton-Watson trees.

We loosely follow the original presentation of the algorithm and deal separately with each case.

Suppose $s_i = w_0$, and consider $(\tau_{i+1}, x_{i+1}, s_{i+1})$; we have $\rho_{i+1}^j = \rho_i^j$ for $j = 1, \dots, n-i-1$, since all the algorithm does is erase x_i (that is, ρ_i^{n-i}), and the claim follows by the induction hypothesis.

If $s_i = w_{>0}$, then the output of the algorithm depends on whether ρ_i^{n-i-1} consists of only $p(x_i)$ or not. If $s_{i+1} = b$ or $s_{i+1} = w_0$, then $\rho_{i+1}^j = \rho_i^j$ for $j = 1, \dots, n-i-2$, which is all that is required. The same is true for $s_{i+1} = w_{>0}$, but we also need to show that $\tau_{i+1}^{r(x_{i+1})}$ is Galton-Watson and independent of the ρ_{i+1}^j 's. The subtree ρ_{i+1}^{n-i-1} is determined by the algorithm as follows. If $\rho_i^{n-i-1} = \{p(x_i)\}$ then it is isomorphic to $\tau_i^{r(x_i)}$ (we are interested in only the case of $r(x_i)$ and $p(x_i)$ being white), which has the required properties by the induction hypothesis. Otherwise we simply have $\rho_{i+1}^{n-i-1} = \rho_i^{n-i-1}$, and we are done.

If $s_i = b$, then we need to deal with a few cases separately. The most straightforward one is that of $s_{i+1} = j$; what the algorithm does in this case is merely erase $\tau_i^{x_i}$: conditions on trees ρ_{i+1}^j are automatic (including the fact that $\rho_{i+1}^{n-i-2} = \{p(x_{i+1})\}$ and $\rho_{i+1}^{n-i-1} = \{x_{i+1}\}$, or we would not have switched to jump state). If $s_{i+1} = w_0$, this may be for one of two reasons:

- x_i has only one right sibling in τ_i , which is a white leaf, and $p(x_i)$ is white (case [b.2]); there is nothing to prove here, since $\rho_{i+1}^j = \rho_i^j$ for $j = 1, \dots, n-i-2$;
- x_i has no right sibling in τ_i , but $p(x_i)$ does, and its next sibling is a white leaf (case [b.3.1]); in this case, while trees ρ_i^j remain unchanged for $j \leq n-i-3$, ρ_{i+1}^{n-i-2} is ρ_i^{n-i-2} with its leftmost branch erased; on the other hand, what we have done is precisely condition such a tree on having a leftmost branch made up of a white leaf, and then remove it, which leaves nothing but a Galton-Watson tree.

The case of $s_{i+1} = b$ presents no added difficulties apart for the need for more casework. One needs to deal separately with cases [b.2], [b.3.1] (same as above, with a weaker condition to check) and [b.1], which is again trivial. Finally, one has to go through essentially the same for $s_{i+1} = w_{>0}$, but with the added requirement to show that $\tau_{i+1}^{r(x_{i+1})}$ is Galton-Watson and independent of the ρ_{i+1}^j 's. This is true in case [b.1] ($\tau_{i+1}^{r(x_{i+1})} = \tau_i^{r(p(x_i))}$), [b.2] ($\tau_{i+1}^{r(x_{i+1})} = \tau_i^{r(r(p(x_i)))}$, where $r(r(p(x_i)))$ does exist in τ_i , or we would not go to state $w_{>0}$), [b.3.1] ($\tau_{i+1}^{r(x_{i+1})} = \tau_i^{r(r(p(p(x_i))))}$).

The very last possibility is for s_i to be j ; if s_{i+1} is j as well, there is hardly anything to prove; all other cases require arguments that are exactly the same as those used for $s_i = b$, and that we shall not repeat.

This concludes the proof by induction. \square

Given Lemma 6.2, Proposition 6.1 is only a matter of computing transition probabilities. We refer the reader to the summary table in the next page.

The purpose of the algorithm was, since the very beginning, to give estimates for the map-distance of vertices from the root; we are now in a position to easily obtain asymptotics for the distance $d^\circ(\text{cut}_n(T_\infty))$. Namely, we have

Proposition 6.3. *Let T_∞ be the geometric Galton-Watson tree conditioned to survive, uniformly bicoloured; then*

$$\lim_{n \rightarrow \infty} \frac{d^\circ(\text{cut}_n(T_\infty))}{n} = 7/9,$$

where the convergence is almost sure. Moreover, for all $\varepsilon > 0$ there exist positive n_ε, C such that, for all $n \geq n_\varepsilon$,

$$\mathbb{P}\left(\left|\frac{d^\circ(\text{cut}_n(T_\infty))}{n} - \frac{7}{9}\right| \geq \varepsilon\right) \leq e^{-Cn}.$$

Proof. We know that, for each $n > 1$, the sequence of random states s_0, \dots, s_{n-2} (from the first $n - 1$ triples that act as input for the algorithm when started on $\text{cut}_n(T_\infty)^\circ$) has the distribution described in Proposition 6.1. Since we have $d^\circ(\text{cut}_n(T_\infty)) = |\{0 \leq i < n | s_i \neq j\}|$, we also have $|d^\circ(\text{cut}_n(T_\infty)) - |\{0 \leq i \leq n - 2 | s_i \neq j\}|| \leq 1$.

On the other hand,

$$\lim_{n \rightarrow \infty} \frac{1}{n} |\{0 \leq i \leq n - 2 | s_i \neq j\}|$$

is a constant by the law of large numbers and is easily computed via the limit distribution for a Markov chain with transition matrix M , which is $\pi = \frac{1}{9}(1, 2, 4, 2)$.

As a consequence, we have

$$\lim_{n \rightarrow \infty} \frac{1}{n} d^\circ(\text{cut}_n(T_\infty)) = \lim_{n \rightarrow \infty} \frac{1}{n} |\{0 \leq i \leq n - 2 | s_i \neq j\}| = 7/9.$$

The second part of the proposition is a direct consequence of classical results of Large Deviation Theory about Markov chains with a finite state space, see for example Chapter 3 of [11]. The statement is true if we substitute $S_n = |\{0 \leq i \leq n - 2 | s_i \neq j\}|$ for $d^\circ(\text{cut}_n(T_\infty))$, because s_0, \dots, s_{n-2} is a (recurrent) Markov chain with finite state space.

The inequality $|d^\circ(\text{cut}_n(T_\infty)) - S_n| \leq 1$ implies that $\mathbb{P}(d^\circ(\text{cut}_n(T_\infty)) = d) \leq \mathbb{P}(|S_n - d| \leq 1)$ for any natural number d , and thus

$$\mathbb{P}\left(\left|\frac{d^\circ(\text{cut}_n(T_\infty))}{n} - \frac{7}{9}\right| \geq \varepsilon\right) \leq \mathbb{P}\left(\left|\frac{S_n}{n} - \frac{7}{9}\right| \geq \varepsilon - \frac{1}{n}\right)$$

which establishes the result for $d^\circ(\text{cut}_n(T_\infty))$. \square

STATE TRANSITIONS

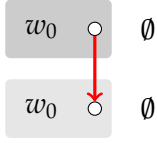
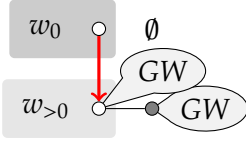
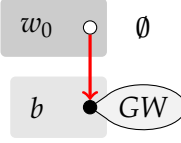
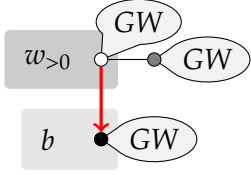
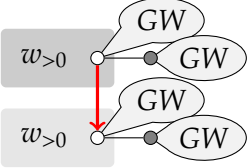
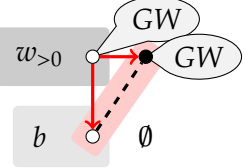
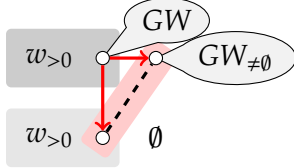
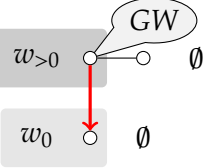
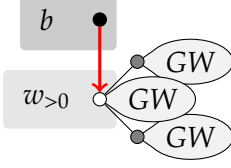
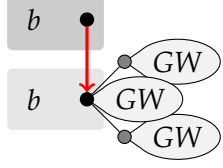
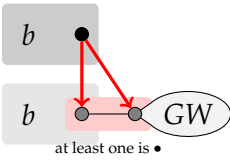
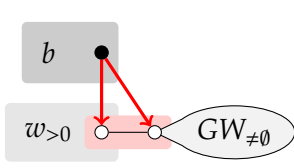
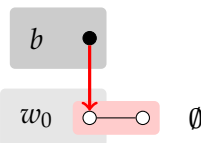
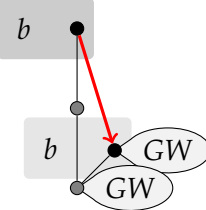
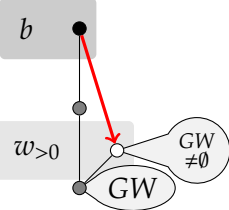
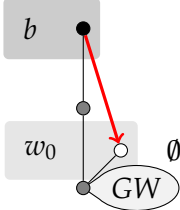
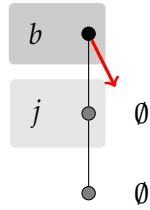
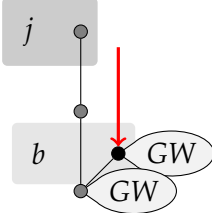
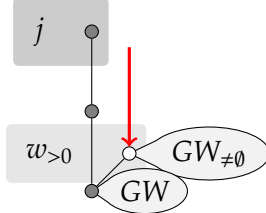
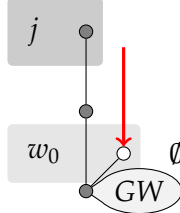
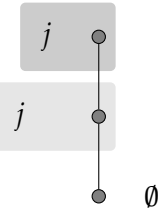
w_0	$\mathbb{P}(\circ)\mu(0)=\frac{1}{4}$  go to w_0	$\mathbb{P}(\circ)(1-\mu(0))=\frac{1}{4}$  go to $w_{>0}$	$\mathbb{P}(\bullet) = \frac{1}{2}$  go to b	$w_{>0}$ $\mathbb{P}(\bullet) = \frac{1}{2}$  go to b	
	$\mathbb{P}(\circ)(1-\mu(0))=\frac{1}{4}$  go to $w_{>0}$	$\mathbb{P}(\circ)\mathbb{P}(\bullet)\mu(0)=\frac{1}{8}$  go to b	$\mathbb{P}(\circ)^2\mu(0)(1-\mu(0))=\frac{1}{16}$  go to $w_{>0}$	$\mathbb{P}(\circ)^2\mu(0)^2=\frac{1}{16}$  go to w_0	
b	$\mathbb{P}(\circ)(1-\mu(0)-\mu(1))=\frac{1}{8}$  go to $w_{>0}$	$\mathbb{P}(\bullet)(1-\mu(0)-\mu(1))=\frac{1}{8}$  go to b	$\mathbb{P}(\circ\bullet,\bullet\circ,\bullet\bullet)\mu(1)=\frac{3}{16}$  go to b	$\mathbb{P}(\circ)^2\mu(1)(1-\mu(0))=\frac{1}{32}$  go to $w_{>0}$	
	$\mathbb{P}(\circ)^2\mu(1)\mu(0)=\frac{1}{32}$  go to w_0	$\mathbb{P}(\bullet)\mu(0)(1-\mu(0))=\frac{1}{8}$  go to b	$\mathbb{P}(\circ)\mu(0)(1-\mu(0))^2=\frac{1}{16}$  go to $w_{>0}$	$\mathbb{P}(\circ)\mu(0)^2(1-\mu(0))=\frac{1}{16}$  go to w_0	$\mu(0)^2 = \frac{1}{4}$  go to j
j	$\mathbb{P}(\bullet)(1-\mu(0))=\frac{1}{4}$  go to b	$\mathbb{P}(\circ)(1-\mu(0))^2=\frac{1}{8}$  go to $w_{>0}$	$\mathbb{P}(\circ)\mu(0)(1-\mu(0))=\frac{1}{8}$  go to w_0	$\mu(0) = \frac{1}{2}$  go to j	

Table 1: Transitions from states w_0 , $w_{>0}$, b and j . Vertices encased within a red region get identified by the algorithm.

7 Final proofs

The final technical steps follow [10], but we need to bypass the rerooting argument exploited there by establishing a more practical control on all map-distances (not only map-distances from the root, as we have done so far). We will start, however, with a proposition dealing only with map-distances from the root, in order to extend the statement as soon as all of the necessary lemmas are in place.

Before we start, let us introduce a little notation: suppose $(x_n)_{n \geq 1}$ is a sequence of real numbers; we write $x_n = \text{oe}(n)$ to mean that $x_n \leq C_1 e^{-C_2 n^a}$ for some positive C_1, C_2, a . Also, in the following proposition and proof, we will write c for the constant $\frac{7}{9}$, which we obtained in Section 6, and d for the usual graph distance on the outerplanar map obtained from a tree via the bijection Ψ , with $d(u)$ being the graph distance of u from the root vertex.

Proposition 7.1. *Let τ_n be a random well bicoloured tree with n vertices, and let $\Delta(\tau_n)$ be its (random) diameter; then for all $\varepsilon > 0$*

$$p_n = \mathbb{P} \left(\exists u \in \tau_n \text{ s.t. } |d(\tau_n, u) - c|u|| \geq \varepsilon \max\{\Delta(\tau_n), \sqrt{n}\} \right) = \text{oe}(n).$$

Proof. For any positive integer k , we call τ_n^k a random well bicoloured tree with n vertices conditioned on having a rightmost branch of length exactly k . We claim that the probability of τ_n having a rightmost branch of length greater than $\varepsilon \sqrt{n}/4$ is $\text{oe}(n)$; we shall need a more general result before the end of this section, and we postpone the proof of this claim, in a stronger form, until the end of the proof (see Lemma 7.2). Given the claim, we have that p_n is $\text{oe}(n)$ if and only if the same is true for

$$\sum_{k=1}^{\frac{\varepsilon}{4}\sqrt{n}} q_{n,k} \mathbb{P} \left(\exists u \in \tau_n^k \text{ s.t. } |d(\tau_n^k, u) - c|u|| \geq \varepsilon \max\{\Delta(\tau_n^k), \sqrt{n}\} \right)$$

where $q_{n,k}$ is the probability that τ_n has a rightmost branch of length exactly k . If we were able to prove that each probability in the sum is actually $\text{oe}(n)$, then we would have $p_n = \text{oe}(n)$ as well.




Figure 6: The tree τ_n^k seen as a forest of k linked rooted plane trees.

Let us consider a single term of the sum. Notice that, for each k , τ_n^k can be seen as a random forest of k ordered trees, with $n - 1$ vertices between them, such that the trees are linked by the roots with a path going from left to right, and a single extra vertex is linked to the root of the rightmost tree; each of the trees is bicoloured (not necessarily *well* bicoloured) with the only condition of having a white root; the rightmost vertex is white.

We label the k (random) trees in the forest τ_n^k , ordered from left to right, $\tau_n^k(1), \dots, \tau_n^k(k)$; if, for each i between 1 and k , we condition $\tau_n^k(i)$ on having a certain number n_i of vertices (with $n_i \geq 1$ and $n_1 + \dots + n_k = n - 1$), then $\tau_n^k(i)$ simply becomes a random plane tree with n_i vertices (uniformly bicoloured but with a white root), which we call θ_{n_i} .

Notice now that, for any vertex u in $\tau_n^k(i)$, $|d(\tau_n^k, u) - d^\circ(\tau_n^k(i), u)| \leq k$; also, the height of u in $\tau_n^k(i)$ differs from the height of u in τ_n^k by at most k , and we have the obvious inequality $\Delta(\tau_n^k(i)) \leq \Delta(\tau_n^k)$ between diameters. As a consequence,

$$\begin{aligned} & \mathbb{P}(\exists u \in \tau_n^k \text{ s.t. } |d(\tau_n^k, u) - c|u|| \geq \varepsilon \max\{\Delta(\tau_n^k), \sqrt{n}\}) \leq \\ & \leq \sum_{i=1}^k \mathbb{P}(\exists u \in \tau_n^k(i) \text{ s.t. } |d^\circ(\tau_n^k(i), u) - c|u|| \geq \varepsilon \max\{\Delta(\tau_n^k(i)), \sqrt{n}\} - 2k) \end{aligned}$$

where in the second expression we still write $|u|$ for the height of the vertex in $\tau_n^k(i)$.

Hence the probability above is no more than

$$\max_{\substack{n_1, \dots, n_k > 0 \\ n_1 + \dots + n_k = n-1}} \sum_{i=1}^k \mathbb{P}(\exists u \in \theta_{n_i} \text{ s.t. } |d^\circ(\theta_{n_i}, u) - c|u|| \geq \varepsilon \max\{\Delta(\theta_{n_i}), \sqrt{n}\} - 2k).$$

Now, we know that $k \leq \frac{\varepsilon}{4} \sqrt{n}$, and therefore we can reduce to evaluating the expression

$$\frac{\varepsilon}{4} \sqrt{n} \max_{0 < m < n} \mathbb{P}(\exists u \in \theta_m \text{ s.t. } |d^\circ(\theta_m, u) - c|u|| \geq \frac{\varepsilon}{2} \max\{\Delta(\theta_m), \sqrt{n}\}).$$

Let θ be a uniformly bicoloured Galton-Watson tree (see Section 5); the probability that θ has m vertices is $2^{-2m-1} \text{Cat}(m-1)$, which is asymptotic to $m^{-\frac{3}{2}}$; we can therefore find a constant C_1 such that, for any m , $2^{-2m-1} \text{Cat}(m-1) \geq C_1 m^{-\frac{3}{2}}$. This guarantees that, for any positive integer m less than n ,

$$\begin{aligned} & \mathbb{P}(\exists u \in \theta_m \text{ s.t. } |d^\circ(\theta_m, u) - c|u|| \geq \frac{\varepsilon}{2} \max\{\Delta(\theta_m), \sqrt{n}\}) \leq \\ & \leq \frac{1}{C_1} m^{\frac{3}{2}} \cdot \mathbb{P}(\exists u \in \theta \text{ s.t. } |d^\circ(\theta, u) - c|u|| \geq \frac{\varepsilon}{2} \max\{\Delta(\theta), \sqrt{n}\}). \end{aligned}$$

If we write q_n for the probability appearing in the above expression, then (since $m \leq n$) proving $q_n = o(n)$ would guarantee that the above – and therefore p_n – is $o(n)$. We shall now turn to the former endeavour.

It is clear that

$$q_n \leq \mathbb{P} \left[\sum_{u \in \theta} \mathbb{1} \left(|d^\circ(\theta, u) - c|u|| \geq \frac{\varepsilon}{2} \max\{\Delta(\theta), \sqrt{n}\} \right) \right]$$

and, grouping vertices together according to height, the latter can be rewritten as

$$\mathbb{P} \left[\sum_{j \geq 1} \sum_{\substack{u \in \theta \\ |u|=j}} \mathbb{1} \left(|d^\circ(\theta, u) - cj| \geq \frac{\varepsilon}{2} \max\{\Delta(\theta), \sqrt{n}\} \right) \right] \leq \sum_{j \geq 1} \mathbb{P} \left[\sum_{\substack{u \in \theta \\ |u|=j}} \mathbb{1} \left(|d^\circ(\theta, u) - cj| \geq \frac{\varepsilon}{2} \max\{j, \sqrt{n}\} \right) \right]$$

where we have used the fact that, for each u in the sum, we have $d(\theta) \geq |u|$.

We can now use Lemma 5.1 to turn the expression above into

$$\sum_{j \geq 1} \mathbb{P} \left(|d^\circ(\text{cut}_j(T_\infty)) - cj| \geq \frac{\varepsilon}{2} \max\{j, \sqrt{n}\} \right)$$

where T_∞ is the critical geometric Galton-Watson tree conditioned to survive, randomly bicoloured, as presented in Section 5.

We split the sum into two parts, which we will deal with separately: the sum for $j \leq n^{1/4}$ and that of the terms with $j > n^{1/4}$.

Suppose that $j \leq n^{1/4}$; clearly, $d^\circ(\text{cut}_j(T_\infty)) \leq j$, and $\max\{j, \sqrt{n}\} = \sqrt{n}$. This gives $|d^\circ(\text{cut}_j(T_\infty)) - cj| \leq d^\circ(\text{cut}_j(T_\infty)) + cj \leq (1+c)j \leq (1+c)n^{1/4}$. Thus it is enough to choose n so that $(1+c)n^{1/4} < \frac{\varepsilon}{2} \sqrt{n}$, that is $n^{1/4} > \frac{2+2c}{\varepsilon}$, to obtain that

$$\sum_{j=1}^{\lfloor n^{1/4} \rfloor} \mathbb{P} \left(|d^\circ(\text{cut}_j(T_\infty)) - cj| \geq \frac{\varepsilon}{2} \max\{j, \sqrt{n}\} \right) = 0.$$

As for the sum with $j > n^{1/4}$, we have

$$\mathbb{P} \left(|d^\circ(\text{cut}_j(T_\infty)) - cj| \geq \frac{\varepsilon}{2} \max\{j, \sqrt{n}\} \right) \leq \mathbb{P} \left(\left| \frac{d^\circ(\text{cut}_j(T_\infty))}{j} - c \right| \geq \frac{\varepsilon}{2} \right)$$

which, by choosing n appropriately according to Proposition 6.3, can be bounded by e^{-Cn} .

This gives, for n suitably big, the bound

$$\sum_{j > n^{1/4}} e^{-Cj} = \text{oe}(n),$$

which is our aim. □

We now prove the claim from the beginning of the proof, in the form of the following Lemma:

Lemma 7.2. *Let τ_n be a random well bicoloured plane tree with n vertices, and let α and ε be positive real numbers; we call a path in τ_n an ancestor path if it is of the form $x_0 \dots x_l$, with $x_i = p(x_{i-1})$ for all i between 1 and l . Then*

$$\mathbb{P}(\tau_n \text{ has an ancestor path of length at least } \alpha n^\varepsilon \text{ entirely made up of white vertices}) = \text{oe}(n);$$

in particular, the probability that τ_n has a rightmost branch of length at least αn^ε is $\text{oe}(n)$.

Proof. We start by showing the last, more specific assertion: that the probability of τ_n having a rightmost branch of length at least αn^ε is $\text{oe}(n)$.

For any $d > 0$, the number of bicoloured trees with n vertices and a rightmost branch of length d (as seen for example in [7]) is

$$2^{n-1-d} \text{Cat}(n-1-d, d) = 2^{n-1-d} \frac{d}{2n-2-d} \binom{2n-2-d}{n-1-d},$$

that is the number of plane trees with n vertices and a rightmost branch of length d (or, equivalently, of sequences of d plane trees with $n-1$ vertices in total, see Figure 6), multiplied by the number of possible bicolourings (2^{n-1-d} , since $d+1$ out of the n vertices belong to the rightmost branch of the tree).

The total number of outerplanar maps with n vertices is asymptotic (up to a multiplicative constant) to $2^{3n} n^{-\frac{3}{2}}$, as can be easily obtained from Stirling estimates for the above formula (see, for detailed analogous computations, [7]); on the other hand,

$$\sum_{\alpha n^\varepsilon < d < n} 2^{n-1-d} \text{Cat}(n-1-d, d) \leq \sum_{\alpha n^\varepsilon < d < n} 2^{n-1-d} \binom{2n-2-d}{n-1-d},$$

which in turn is less than

$$2^{3n} p(n) \sum_{\alpha n^\varepsilon < d < n} 2^{-2d},$$

where $p(n)$ a polynomial in n . As a consequence, the above expression divided by the total number of outerplanar maps with n vertices is $\text{oe}(n)$.

It is now very easy to extend the result to general ancestor paths. Since the probability that τ_n has a rightmost branch of length at least $\frac{\alpha}{2} n^\varepsilon$ we have shown to be $\text{oe}(n)$, we may assume τ_n is conditioned on having a rightmost branch of length less than $\frac{\alpha}{2} n^\varepsilon$. For each vertex v of height at least αn^ε in τ_n consider the path $P_v = vp(v) \dots p^i(v) \dots p^{\lceil \alpha n^\varepsilon \rceil}(v)$. No more than $\frac{1}{2} \alpha n^\varepsilon$ of its vertices belong to the rightmost branch of τ_n , and thus the probability of P_v being entirely white is at most $2^{-\frac{\alpha}{2} n^\varepsilon}$. Hence

$$\mathbb{P} \left(\begin{array}{l} \tau_n \text{ has an ancestor path of length } \alpha n^\varepsilon \text{ entirely made up of white} \\ \text{vertices and the rightmost branch of } \tau_n \text{ is shorter than } \frac{\alpha}{2} n^\varepsilon \end{array} \right) < n \cdot 2^{-\frac{\alpha}{2} n^\varepsilon}$$

which is $\text{oe}(n)$ as wanted. □

Proposition 7.1 is a substantial step toward being able to bound the Gromov-Hausdorff distance between a (rescaled) tree and its corresponding planar map, but dealing with distances from the root is not enough: we need a way to derive results of the same kind about distances between generic vertices.

To this end, we will rewrite the map-distance between two vertices in terms of the distances between each vertex and the root. This is easily done when the vertices in question are related, and the one of smaller height is black; this basic case we will use as a stepping stone, together with Lemma 7.2, to establish the required general result.


Lemma 7.3. *Let τ be a bicoloured tree and v a vertex in τ ; let w be a black ancestor of v in τ , the root of τ ; call d_M the map-distance on τ and (as in Section 3) write $d_M(u)$ for $d_M(u, \emptyset)$. Then $|d_M(v, w) - d_M(v) + d_M(w)| \leq 2$.*

Proof. Consider a map-geodesic from v to the root; if this path goes through w , then $d_M(v) = d_M(v, w) + d_M(w)$. If it does not, then at some point it jumps from a descendant s of w onto a target t ‘below’ w , which must be the child of an ancestor of w , thus also the target of w . We have

$$d_M(v, s) \leq d_M(v, w) \leq d_M(v, s) + 2$$

$$d_M(v) = d_M(v, s) + 1 + d_M(t)$$

$$d_M(t) \leq d_M(w) \leq d_M(t) + 1$$



hence $|d_M(v) - d_M(v, w) - d_M(w)| \leq 2$ as wanted. \square

Here is a general statement analogous to Proposition 7.1, where we write $d_T(u, v)$ for the distance of two vertices *in the tree*, and d_M for the map distance (again, with $d_M(u)$ being the map distance from the root).

Corollary 7.4. *Let τ_n be a random well bicoloured tree with n vertices, and let $\Delta(\tau_n)$ be its (random) diameter; then for all $\varepsilon > 0$*

$$\mathbb{P}(\exists u, v \in \tau_n \text{ s.t. } |d_M(u, v) - cd_T(u, v)| \geq \varepsilon \max\{\Delta(\tau_n), \sqrt{n}\}) = o(n).$$

Proof. Thanks to Lemma 7.2 (by choosing $\alpha = \frac{\varepsilon}{8}$) we may restrict ourselves to the event of τ_n having no white ancestor path of length $\frac{\varepsilon}{8} \sqrt{n}$ or greater.

Consider, given u and v vertices of τ_n , their first common ancestor w . Either $|w| > \frac{\varepsilon}{8} \sqrt{n}$, in which case there is a *black* ancestor z of w such that $d_T(w, z) \leq \frac{\varepsilon}{8} \sqrt{n}$ (otherwise there would be a long white ancestor path), or $|w| < \frac{\varepsilon}{8} \sqrt{n}$, in which case we just set z to be the root of τ_n .

Suppose without loss of generality that u lies to the left of v and let w' be the child of w that is also an ancestor of v ; then a map-geodesic from u to v goes through at least one of w and w' , thanks to an argument very similar to that employed in Proposition 3.1. This yields that $|d_M(u, v) - d_M(u, w) - d_M(w, v)| \leq 2$ (this is trivial if the map-geodesic passes through w ; if it goes through w' then we have $|d_M(u, v) - d_M(u, w) - d_M(w, v)| = |d_M(u, w') + d_M(w', v) - d_M(u, w) - d_M(w, v)| \leq |d_M(u, w') - d_M(u, w)| + |d_M(w', v) - d_M(w, v)| \leq 2$).

Also notice that we have $|d_M(u, w) - d_M(u, z)| \leq d_M(w, z) \leq \frac{\varepsilon}{8} \sqrt{n}$, and the same inequality is true if we substitute v for u .

Now, if z is black, then $|d_M(u, z) - d_M(u) + d_M(z)| \leq 2$ and $|d_M(v, z) - d_M(v) + d_M(z)| \leq 2$ by Lemma 7.3; otherwise z is the root of τ_n , and the same assertions are trivial (since $d_M(u, z) = d_M(u)$ and $d_M(z) = 0$).

All of the above observations combined yield

$$\begin{aligned} |d_M(u, v) - cd_T(u, v)| &= |d_M(u, v) - c(|u| + |v| - 2|w|)| \leq \\ &\leq \frac{\varepsilon}{2} \sqrt{n} + 6 + |d_M(u) - c|u|| + |d_M(v) - c|v|| + 2|d_M(z) - c|z||. \end{aligned}$$

Hence

$$\begin{aligned} &\mathbb{P}(\exists u, v \in \tau_n \text{ s.t. } |d_M(u, v) - cd_T(u, v)| \geq \varepsilon \max\{\Delta(\tau_n), \sqrt{n}\}) \leq \\ &\mathbb{P}(\exists u, v, z \in \tau_n \text{ s.t. } \frac{\varepsilon}{2} \sqrt{n} + 6 + |d_M(u) - c|u|| + |d_M(v) - c|v|| + 2|d_M(z) - c|z|| \geq \varepsilon \max\{\Delta(\tau_n), \sqrt{n}\}) \leq \\ &\mathbb{P}(\exists u, v, z \in \tau_n \text{ s.t. } |d_M(u) - c|u|| + |d_M(v) - c|v|| + 2|d_M(z) - c|z|| \geq \frac{\varepsilon}{4} \max\{\Delta(\tau_n), \sqrt{n}\}) \leq \\ &\mathbb{P}(\exists u \in \tau_n \text{ s.t. } |d_M(u) - c|u|| \geq \frac{\varepsilon}{16} \max\{\Delta(\tau_n), \sqrt{n}\}) \end{aligned}$$

which we show to be $oe(n)$ by invoking Proposition 7.1. \square

The time has come for the proof of our main theorem, which is now quite straightforward; we restate it here:

Theorem 1.1. *Let \mathbf{M}_n be a random uniform rooted simple outerplanar map with n vertices, and denote by d_{gr} the graph distance on the set of its vertices $V(\mathbf{M}_n)$. We have the following convergence in distribution for the Gromov-Hausdorff topology:*

$$\left(V(\mathbf{M}_n), \frac{d_{gr}}{\sqrt{n}} \right) \xrightarrow[n \rightarrow \infty]{(d)} \frac{7\sqrt{2}}{9} \cdot (\mathcal{T}_e, d),$$

where (\mathcal{T}_e, d) is the Brownian CRT of Aldous. We adopt here the normalisation of Le Gall [14] by considering \mathcal{T}_e as constructed from a normalised Brownian excursion.

Proof. Corollary 7.4 yields that, given a random well bicoloured tree with n vertices τ_n ,

$$\mathbb{P}(d_{GH}(\Psi(\tau_n), c\tau_n) \geq \varepsilon \max\{\sqrt{n}, \Delta(\tau_n)\}) = oe(n)$$

where $\Psi(\tau_n)$ in this context is seen as the metric space made up of the vertices of τ equipped with the map-distance, and $c\tau_n$ is the set of vertices in τ_n , equipped with the graph distance of τ_n rescaled by a factor $c = \frac{7}{9}$.

This is because

$$d_{GH}(\Psi(\tau_n), c\tau_n) \leq \frac{1}{2} \left(\sup_{u, v \in \tau_n} |d_M(u, v) - cd_T(u, v)| \right),$$

as seen by considering the trivial correspondence between the vertices of τ_n and those of $\Psi(\tau_n)$.

Thus we have established that

$$\lim_{n \rightarrow \infty} d_{GH} \left(\frac{\Psi(\tau_n)}{\max\{\sqrt{n}, \Delta(\tau_n)\}}, \frac{c\tau_n}{\max\{\sqrt{n}, \Delta(\tau_n)\}} \right) = 0$$

in probability, where τ_n is a random rooted well bicoloured plane tree with n vertices.

We now claim that

$$\lim_{n \rightarrow \infty} \left(\tau_n, \frac{d_{gr}}{\sqrt{2n}} \right) = (\mathcal{T}_e, d)$$

in distribution for the Gromov-Hausdorff distance, with (\mathcal{T}_e, d) being the CRT.

This result is a consequence of a famous theorem of Aldous [5] and would be immediate if τ_n were replaced by a uniform plane tree with n vertices or, equivalently, by a critical geometric Galton-Watson tree conditioned on having n vertices. Even though this is not the case, τ_n is not very far from the latter: indeed, we saw in Lemma 7.2 that the length L_n of the rightmost branch of τ_n remains tight (it even converges in distribution) as $n \rightarrow \infty$, and furthermore that, conditionally on $L_n = k$, the k subtrees grafted onto the rightmost branch form a forest $\tau_n^k(1), \tau_n^k(2), \dots, \tau_n^k(k)$ whose total number of vertices is $n - 1$. It is known that such a forest has, as $n \rightarrow \infty$, a unique macroscopic tree $\tilde{\tau}_n$ of size $s_n = n - o(n)$, which is uniformly distributed over all plane trees of size s_n . The scaling limit of τ_n is thus the same as that of $\tilde{\tau}_n$, which is that of a random uniform plane tree of size n . See Section 3.3 of [17] for details.

But then the random variable

$$\frac{\max\{\sqrt{n}, \Delta(\tau_n)\}}{\sqrt{n}}$$

also converges in distribution, and to an almost surely positive random variable. That is, if we multiply by $\frac{\max\{\sqrt{n}, \Delta(\tau_n)\}}{\sqrt{n}}$ we can in fact deduce that

$$\lim_{n \rightarrow \infty} d_{GH} \left(\frac{\Psi(\tau_n)}{\sqrt{n}}, \frac{c\tau_n}{\sqrt{n}} \right) = 0.$$

Now remember that, thanks to Theorem 2.1, $\Psi(\tau_n)$ is a random rooted simple out-erplanar map with n vertices, that is it has the same distribution as \mathbf{M}_n . This finally gives

$$\lim_{n \rightarrow \infty} \left(\mathbf{M}_n, \frac{d_{gr}}{\sqrt{n}} \right) = \lim_{n \rightarrow \infty} \left(\tau_n, \frac{7d_{gr}}{9\sqrt{n}} \right) = \left(\mathcal{T}_e, \frac{7\sqrt{2}}{9}d \right).$$

□

Acknowledgements

I would like to thank N. Curien for his numerous suggestions and improvements to the readability of this paper; J-F. Le Gall and F. Flandoli for their supervision and support; finally, D. Lombardo for the many long and in-depth discussions about the subject matter.

References

- [1] ADDARIO-BERRY, L., BROUTIN, N., AND GOLDSCHMIDT, C. The continuum limit of critical random graphs. *Probab. Theory Relat. Fields* 152, 3-4 (2012), 367–406.
- [2] ALBENQUE, M., AND MARCKERT, J.-F. Some families of increasing planar maps. *Electron. J. Probab.* 13 (2008), no. 56, 1624–1671.
- [3] ALDOUS, D. The continuum random tree. I. *Ann. Probab.* 19, 1 (1991), 1–28.
- [4] ALDOUS, D. *The Continuum random tree II: an overview*. London Mathematical Society Lecture Note Series. Cambridge University Press, 1991.
- [5] ALDOUS, D. The continuum random tree. III. *Ann. Probab.* 21, 1 (1993), 248–289.
- [6] BETTINELLI, J. Scaling Limit of Random Planar Quadrangulations with a Boundary. *Annales de l’institut Henri Poincaré* (2011), to appear.
- [7] BONICHON, N., GAVOILLE, C., AND HANUSSE, N. Canonical decomposition of outerplanar maps and application to enumeration, coding and generation. *J. Graph Algorithms Appl.* 9, 2 (2005), 185–204.
- [8] BOUTTIER, J., AND GUITTER, E. Distance statistics in quadrangulations with a boundary, or with a self-avoiding loop. *J. Phys. A* 42, 46 (2009), 465208, 44.
- [9] CHARTRAND, G., AND HARARY, F. Planar permutation graphs. *Annales de l’institut Henri Poincaré (B) Probabilités et Statistiques* 3, 4 (1967), 433–438.
- [10] CURIEN, N., HAAS, B., AND KORTCHEMSKI, I. The CRT is the scaling limit of random dissections. *Random Struct. Algorithms*, to appear.
- [11] DEMBO, A., AND ZEITOUNI, O. *Large deviations techniques and applications*, vol. 38 of *Stochastic Modelling and Applied Probability*. Springer-Verlag, Berlin, 2010.
- [12] HAAS, B., AND MIERMONT, G. Scaling limits of Markov branching trees with applications to Galton-Watson and random unordered trees. *Ann. Probab.* 40, 6 (2012), 2589–2666.
- [13] JANSON, S., AND ORN STEFANSSON, S. Scaling limits of random planar maps with a unique large face. *Ann. Probab.*, to appear.

- [14] LE GALL, J.-F. Random trees and applications. *Probab. Surv.* 2 (2005), 245–311.
- [15] LE GALL, J.-F. Uniqueness and universality of the Brownian map. *Ann. Probab.* 41, 4 (2013), 2880–2960.
- [16] LYONS, R., AND PERES, Y. Probability on trees and networks, 2005.
- [17] MARCKERT, J.-F., AND PANHOLZER, A. Noncrossing trees are almost conditioned Galton-Watson trees. *Random Struct. Algorithms* 20, 1 (2002), 115–125.
- [18] MIERMONT, G. The Brownian map is the scaling limit of uniform random plane quadrangulations. *Acta Math.* 210, 2 (2013), 319–401.
- [19] PANAGIOTOU, K., STUFLER, B., AND WELLER, K. Scaling limits of random graphs from subcritical classes, in preparation.
- [20] SYSŁO, M. M. Characterizations of outerplanar graphs. *Discrete Math.* 26, 1 (1979), 47–53.